Iterative importance sampling algorithms for parameter estimation problems

Matthias Morzfeld[1], Marcus S. Day[2], Ray W. Grout[3], George Shu Heng Pau[2], Stefan A. Finsterle[2], John B. Bell[2]

[1] Department of Mathematics, University of Arizona;
[2] Lawrence Berkeley National Laboratory;
[3] National Renewable Energy Laboratory;

arXiv:1608.01958v1 [math.NA] 5 Aug 2016

Abstract

In parameter estimation problems one approximates a posterior distribution over uncertain parameters defined jointly by a prior distribution, a numerical model, and noisy data. Typically, Markov Chain Monte Carlo (MCMC) is used for the numerical solution of such problems. An alternative to MCMC is importance sampling, where one draws samples from a proposal distribution, and attaches weights to each sample to account for the fact that the proposal distribution is not the posterior distribution. Drawing samples and computing their weights can be done independently, which leads to massively parallel algorithms that exhibit near perfect scaling with the number of cores. These advantages however are conditional on ones ability to construct an appropriate proposal distribution. Here we present iterative algorithms that refine the mean and covariance matrix of Gaussian or multivariate t-proposal distributions. Such ideas have been discussed in the literature before. We revisit and revise these algorithms and investigate their applicability by testing their performance in two realistic test problems, one in subsurface flow, and one in combustion modeling. In our implementation, the algorithms leverage massively parallel computers, and we also discuss how to use coarse MCMC simulations or optimization in connection with Gaussian mixture models to initialize the iterations.

1 Introduction

Predicting the behavior of complex physical systems is a key requirement in science and engineering. One approach to predicting the behavior of such systems is through high fidelity simulation. However, for many systems, model uncertainties limit predictive capability. The idea in parameter estimation is to use experimental data to reduce model uncertainties and, therefore, improve predictions of overall system behavior as follows. One represents uncertainties in the model by assuming that model parameters are random variables with given "prior" distributions, and one represents mismatch between model and data by a "likelihood". By Bayes' rule, the prior and likelihood define a posterior distribution that describes how well system behavior can be predicted based on available data – which aspects are tightly bounded and which are less precise. The posterior distribution is thus regarded as the solution of a parameter estimation problem [17, 75, 76].

Numerical solution of parameter estimation problems typically relies on Monte Carlo (MC) sampling because it provides full insight into the characterization of uncertainty of a complex system. The key issue with a Bayesian approach is that MC techniques may not be robustly effective for complex, high-dimensional problems [16, 48, 60]. For example, particle filter methods break down in high dimensions because events that have a high probability with respect to the prior may have a low probability with respect to the posterior distribution [5, 8, 18, 70–72]. Metropolis and heat bath Markov Chain Monte Carlo (MCMC) algorithms may converge slowly on problems that are poorly conditioned [35, 39].

The central theme of this paper is to present a MC sampling method that is efficient in parameter estimation problems for complex systems, and well suited to extreme-scale computer architectures. We leverage parallelism by using importance sampling, rather than MCMC, so that samples can be generated and weighted in parallel. We make importance sampling robustly applicable to parameter estimation problems with non-convex posterior distributions by iteratively updating the proposal distribution (see section 2.5 for a more precise definition of the problems we want to tackle). The idea of using an iteration to improve efficiencies of proposal distributions has been discussed before, see, e.g., [15, 57, 58], and the literature review below. Our emphasis is on investigating whether iterative methods can indeed be effective in realistic problems. Specifically, we apply iterative sampling methods in two test cases: one in estimation of permeabilities in a subsurface flow problem, and another in estimation of parameters of a combustion model.

2 Background and problem formulation

2.1 Parameter estimation problem formulation

We consider a set of parameters θ of a numerical model $\mathcal{M}(\theta)$. Here θ is an n_θ-dimensional vector, and the model $\mathcal{M} : R^{n_\theta} \to R^{n_z}$ maps the parameter vector to an n_z dimensional vector of outputs z. For example, θ can be a collection of reaction rate parameters and third-body coefficients that define a combustion model, or θ can define a permeability field of a subsurface flow problem (see sections 4 and 5 below). Uncertainty in the model is represented by uncertainty in the parameters, i.e., θ is a random vector with prior probability distribution p_θ. The outputs of the model can be measured, and additional uncertainty in the approximation of the measurements are represented by a random variable v with given distribution p_v. We assume that this uncertainty is additive,

$$z = \mathcal{M}(\theta) + v, \tag{1}$$

however this assumption is not central for our approach. The above equation defines the likelihood $p(z|\theta) = p_v(z - \mathcal{M}(\theta))$. By Bayes' rule, prior and likelihood define the posterior distribution

$$p(\theta|z) \propto p_\theta(\theta)p(z|\theta),$$

which describes the probability of the parameters θ given the data z. The posterior distribution allows for a complete description of the uncertainty in the parameters given the data. In particular, one can compute least-squares optimal estimates by computing the posterior mean as well as "error bars" based on posterior covariances.

2.2 Review of numerical methods for parameter estimation

There is a large literature on Monte Carlo methods. Here we review some techniques that are derived for, or used in, problems where simulations by the model $\mathcal{M}$ are computationally expensive.

If the prior and the errors v are Gaussian, and if, in addition, the model $\mathcal{M}$ is linear, then the posterior is also Gaussian. In this case, it is sufficient to compute its mean and covariance. This can be done, for example, by minimizing the negative logarithm of the posterior distribution

$$F(\theta) = -\log p(\theta|z) + C = -\log p_\theta(\theta) - \log p(z|\theta) + C.$$

The minimizer of F is the posterior mean, and the Hessian, $H_{ij} = \partial^2 F/\partial\theta_i\partial\theta_j$, $i, j = 1, \ldots, n_\theta$ is the inverse of the posterior covariance matrix [76]. Linear algorithms have been made efficient for large-scale linear problems [13, 29].

When the model is nonlinear, the posterior distribution is no longer Gaussian, even if the prior distribution for the parameters and the model errors, v, remain Gaussian. In this case, minimizing the negative log-posterior $F(\theta)$ results in the most likely state given the data, which is called "posterior mode". However, there may be more than one posterior mode (local maxima of the posterior distribution), in which case it is difficult to argue that a (local) posterior mode is a meaningful solution of the parameter estimation problem. For that reason, Markov Chain Monte Carlo (MCMC) methods are routinely used for solving nonlinear parameter estimation problems. The basic idea is to pick a proposal distribution to propose a "move", i.e., a new set of parameters. This move is accepted or rejected based on a suitable accept/reject criterion [41]. The chain then consists of a number of correlated samples and converges in the sense that averages over the samples converge to expected values with respect to the posterior distribution as the ensemble size goes to infinity. Thus, MCMC provides a more complete description of the uncertainty than computing modes of the posterior distribution. In practice, the beginning of an MCMC chain is often disregarded as a "burn-in" period. In addition, the correlation "time" of the chain should not be too large; an approximation by a finite chain may not be appropriate otherwise. The various MCMC algorithms in the literature differ in their proposal distributions and approaches to reduce burn-in and auto-correlation time.

There are several "classical" MCMC algorithms, e.g., Metropolis-Hastings [41], and many of them are known to be slow in high-dimensional and highly nonlinear problems, see, e.g., [35,39]. Recent MCMC methods that address these problems include affine invariant ensemble samplers [30,35], adaptive MCMC [37,38], and differential evolution MCMC [9,10]. Another class of modern MCMC algorithms relies on solving the negative log-posterior optimization problem, $\min_\theta F(\theta)$, repeatedly. Specifically, stochastic Newton MCMC [50, 55], makes use of adjoints during the optimization to speed up gradient and Hessian computations. A related method is "randomize-then-optimize" (RTO) [4], where a stochastic version of the negative log-posterior optimization problem is solved to produce samples. Both, RTO and stochastic Newton are efficient if the problem is linear, however they may struggle on "highly nonlinear problems" with multiple modes.

There is also a large literature on MCMC methods in function spaces, see, e.g., [7, 21, 68, 75]. These algorithms apply to parameter estimation problems where the model $\mathcal{M}$ is a PDE. The motivation for this class of algorithms is that the computational requirements of MCMC algorithms, i.e., burn-in or autocorrelation times, should not increase when the mesh is refined. In fact, these properties should be invariant under mesh-refinement, once grid convergence is reached. For that reason, function space MCMC methods are also often called "mesh-independent" or "dimension-independent". Their theory is best understood for parameter estimation problems connected with elliptic PDE.

Computational requirements of MCMC can also be reduced by solving a related, but easier problem. This idea is implemented by reduced-order modeling: rather than using the model $\mathcal{M}$ for MCMC, one replaces the model by a simpler, "reduced-order" version of it. Since the reduced-order model is simpler, it is computationally less expensive, and using it therefore speeds up MCMC. In this context, one can reduce the dimension of the parameter vector θ, the dimension of the state of the underlying numerical model, $\mathcal{M}$, or both, see, e.g., [22, 23, 45, 73]. Reducing the state dimension is usually achieved by finding a simplified approximation to the mathematical model $\mathcal{M}$. This reduced model however may depend on the same number of parameters as the "full" model. To reduce this number of parameters one can determine a low-dimensional subspace of parameter combinations which are constrained by data, as in the likelihood-informed methods [22, 23]. In problems where the parameters are spatial fields, e.g., permeability fields in subsurface flow problems, one can use principal component analysis/Karhunen-Loeve expansions [53], to define this field by a small number of dominant modes (see also section 4, where Kriging is used to

reduce the parameter dimension). In either case, MCMC by reduced-order models requires that the problem indeed has the assumed low-dimensional structure, or else the posterior distribution defined by the reduced-order model may be very different from the posterior distribution defined by $\mathcal{M}$. Such ideas were made precise in the context of polynomial chaos expansions in [47]. However, it is important to emphasize that reduced-order modeling approaches are only applicable if the model $\mathcal{M}$ indeed has a low-dimensional representation that can be discovered by reduced order modeling techniques.

Besides MCMC, researchers have recently explored the use of optimal transport maps to solve parameter estimation problems [54, 59, 66, 78]. The idea is to construct a map that converts prior samples to posterior samples. More generally, the map should transform a reference distribution to the posterior distribution [16]. However, constructing an optimal map requires solving an infinite dimensional optimization problem, which is expensive. Alternatively, crude approximations to an optimal map may be more readily computable and could be used to "speed up" MCMC. Work in this direction is underway [61, 62].

2.3 Overview of importance sampling

An alternative to MCMC is importance sampling. Importance sampling, just as MCMC, requires a proposal distribution q. However, rather than using the proposal distribution for defining a Markov chain, one draws samples from it and attaches to each sample a weight

$$w(\theta) \propto \frac{p(\theta|z)}{q(\theta)}.$$

The weighted samples form an empirical estimate of the posterior distribution in the sense that weighted averages over the samples converge to expected values as the ensemble size goes to infinity, see, e.g., [17,60]. Moreover, the posterior distribution and the proposal distribution need be known up to a multiplicative constant (as in MCMC), because the weights can be self-normalized such that their sum is one, i.e., $w_i \leftarrow w_i / \sum_{j=1}^{N_e} w_j$. Advantages of importance sampling compared to MCMC are that there is no burn-in time, and the weighted samples are independent. The latter means in particular that the samples and weights can be computed independently, which can lead to enormous computational advantages on massively parallel computer architectures.

However, the efficiency of importance sampling hinges on constructing a suitable proposal distribution q. The reason is that, even though all samples are independent, each sample carries a different weight and, therefore, the samples are not equally important. For an effective importance sampling method all samples should be similarly important, i.e., the variance of the weights should be small. Indeed, a heuristic "effective sample size" can be defined [2, 6, 28, 34, 46, 79] based on the relative variance $var(w)/E(w)^2$:

$$N_{\text{eff}} = \frac{N}{R}, \quad R = \frac{var(w)}{E(w)^2} + 1 = \frac{E(w^2)}{E(w)^2}, \tag{2}$$

The effective sample size describes the size of an unweighted ensemble of size N_{eff} that is equivalent to a weighted ensemble of size N. In particular, suppose that one weight is close to one, which implies that all other self-normalized weights are close to zero. In this case, $R \approx N$, so that the effective sample size is one. If this happens, the sampling method is said to have "collapsed". In contrast, if the proposal distribution is equal to the posterior distribution, the variance of the weights is zero, so that R is unity. For effective sampling, one needs an R that is "not too large", where what large is depends mainly on how much computing power one is willing to spend.

The collapse of sequential Monte Carlo methods for time-dependent problems, called particle filters [28], has been studied extensively in the meteorological literature, and it was shown that, for a certain class of problems, $N_{\text{eff}} \to 1$ as the "effective dimension" of the problem increases [5, 8, 16, 18, 70–72]. The effective dimension can depend on the dimension of the data vector z, the dimension of the parameter vector θ, as well as the distributions of errors in the prior and likelihood and may increase when the dimension of the parameter vector θ increases, however, at a slower rate. Moreover, recent work suggests that particle filters may be efficient even if the effective dimension is large, provided that the filtering problem has a certain "sparse" structure [65].

2.4 Review of implicit sampling

Implicit sampling is an importance sampling method that has been applied to time-dependent "filtering" problems [3, 19, 20, 52] and, more recently, to a relatively simple parameter estimation test problem [53]. The idea is to construct a proposal distribution that is centered at the posterior mode. The posterior mode is the minimizer of

$$\mu = \arg\min_{\theta} F(\theta).$$

Implicit sampling in its simplest implementation amounts to using a Gaussian proposal distribution whose mean is the posterior mean μ, and whose covariance is the inverse of the Hessian of F evaluated at the minimizer. The weight of each Gaussian sample is

$$w(\theta) \propto \exp\left(F_o(\theta) - F(\theta)\right),$$

where F_o is the Taylor expansion of F to second order, $F_0 = \phi + 0.5(\theta - \mu)^T H^{-1}(\theta - \mu)$. Here, ϕ is the minimum of F and superscript T means transpose. This method is also discussed in [60], and it is suggested to use a multivariate t-distribution, rather than a Gaussian, to avoid using a proposal distribution whose tails decay too quickly. Other implementations of implicit sampling, e.g., by random maps, are discussed in [19, 52]. In [34], implicit sampling by random maps or by Gaussian approximation are analyzed in a small noise limit (see [79] for small noise analysis of other sampling methods). The result is that implicit sampling is efficient for nearly Gaussian problems, for which $F(\theta) = F_o(\theta) + O(\varepsilon)$, i.e., the posterior is Gaussian with a "small" perturbation. Here, "efficient" means that $R = 1 + K(n_\theta)O(\varepsilon)$, where n_θ is the parameter dimension. In particular, small noise analysis shows that implicit sampling can be "quite useful in dimensions of practical interest" [34]. Moreover, a random map implementation is more efficient than the implementation by Gaussian approximation, in the sense that R is smaller. However, as the dimension increases, these advantages disappear.

2.5 Motivation and problem definition

Our motivation for using iterative sampling algorithms comes from encountering posterior distributions that are difficult to deal with by the MCMC/importance sampling technology reviewed above. Specifically, we need algorithms that work efficiently when the negative logarithm of the posterior, $F(\theta)$, is not convex and has multiple minima that may not be well separated. We wish to exploit massively parallel computer architectures and, therefore, wish to use importance sampling rather than MCMC. More specifically, we wish to create importance sampling algorithms that are efficient under the following conditions.

1. The priors are "broad", i.e., using the prior distribution as proposal distribution leads to the collapse of importance sampling.

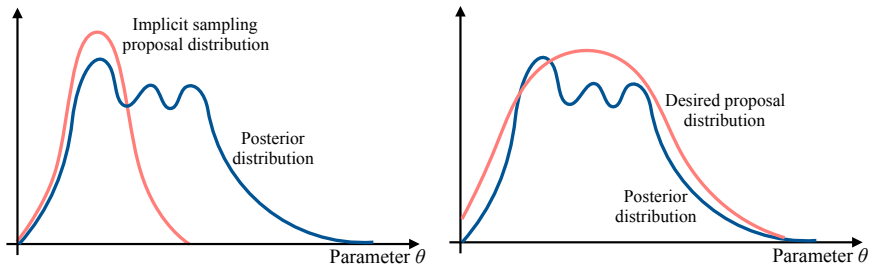

Figure 1: Left: illustration of why implicit sampling may perform poorly on a problem with multiple modes. The proposal density is centered at one of the modes and defines a local, but not global approximation of the posterior distribution. Right: illustration of an effective proposal distribution.

2. We assume that we can compute first derivatives of the negative log-posterior distribution $F(\theta)$, but not second derivatives. This situation is usually encountered in complex computational models, e.g., if adjoints are available, but second-order adjoints are not.

3. We assume that the negative log-posterior distribution $F(\theta)$ is not convex and has multiple minima that may not be well separated.

4. We assume that the effective dimension, n_θ of the problem is moderate, perhaps tens, rather than hundreds or even thousands, and that the model $\mathcal{M}$ can be evaluated by a powerful desktop machine. This will allow us to perform effective parameter estimation on a massively parallel machine.

Under the above conditions, many MCMC methods may converge slowly, due to multiple modes that can trap the chains. Implicit sampling can be expected to perform poorly, since the method relies on the assumption that the posterior distribution has only one significant mode. More specifically, implicit sampling would perform poorly if the posterior distribution is characterized by a plateau of high probability with several small "dents", as illustrated in the left panel of figure 1. In this case, implicit sampling would construct a proposal distribution that is centered at a local mode, with the covariance reflecting local curvature of $F(\theta)$. As a result, the proposal distribution decays too quickly to be useful for parameters further away from the local mode. One may try to apply implicit sampling at each of the modes and "mix" the results by a mixture model, and we present an example of this in section 5. However, we find that a proposal distribution with multiple mixture components can be inefficient because the local curvature may not be informative about the global "structure" of the function. Indeed, the goal of the iterative sampling methods we discuss is to construct a proposal distribution that captures global characteristics of non-convex log-posteriors, $F(\theta)$, as illustrated in the right panel of figure 1.

3 Iterative importance sampling

We wish to iteratively construct an effective proposal distribution from a small number of samples that were computed by other means, e.g., a "short" MCMC simulation, or by a mixture model (see section 3.3 for initialization strategies). Specifically, the initial samples are used to define a proposal distribution q^0, which we refine by iteration as follows. We use N_e^0 weighted samples $\{\theta_j^1, w_j^1\}$, $j = 1, \ldots, N_e^0$, we obtained by importance sampling with q^0 to define a new proposal distribution q^1, and repeat. During the iteration, we monitor the quality measure R in equation (2), and stop

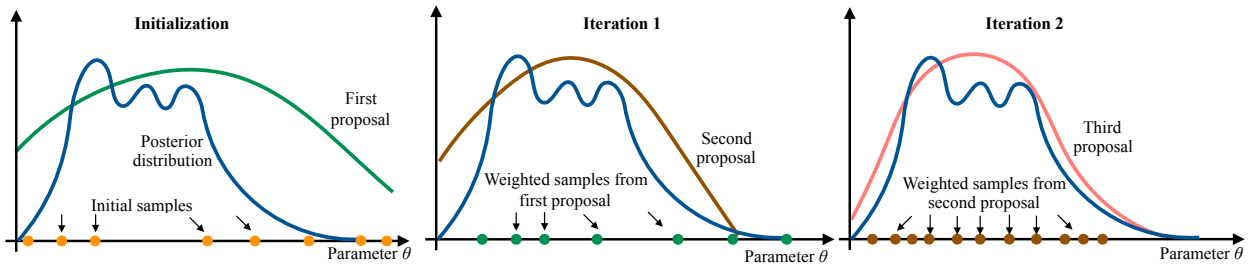

Figure 2: Illustration of iterative construction of a proposal distribution.

Algorithm 1 Iterative importance sampling algorithm

Given: an unweighted ensemble $\left\{\theta_j^0\right\}$, $j = 1, \ldots, N_e$
Construct a proposal distribution q^0 from the ensemble.
Set $k = 0$, and define a tolerance tol
while no convergence reached **do**
  Importance sampling with proposal distribution q^k to get $\left\{\theta_j^{k+1}, w_j^{k+1}\right\}$, $j = 1, \ldots, N_e^k$
  Compute R^{k+1} by equation (2) for this set of weighted samples
  if $R^{k+1}/R^k <$ tol **then**
    Convergence is reached
  else
    Construct proposal distribution q^{k+1} from $\left\{\theta_j^{k+1}, w_j^{k+1}\right\}$, $j = 1, \ldots, N_e^k$
  end if
  Set $k \leftarrow k + 1$
end while

the iteration when $|R_{k+1} - R_k|$ drops below a specified tolerance. The algorithm is summarized in pseudo-code in algorithm 1, and a successful iteration is illustrated in figure 2. The rationale for an iterative approach is to discover, via iteration, the global structure of the posterior distribution without being stuck in local minima as illustrated by figure 1. The success, convergence and efficiency of iterative importance sampling algorithms (ISA) depend on initialization, the proposal distributions used at each step, the ensemble size at each step, and, of course, on the parameter estimation problem itself.

3.1 Discussion of convergence of iterative importance sampling

The iteration we describe converges in the following sense. Suppose that, at each iteration, we use a proposal distribution whose support includes the support of the posterior distribution. Then, at each iteration, weighted samples from the proposal distribution q^k generate an "empirical estimate" of the proposal distribution that converges in the sense that weighted averages converge to expected values with respect to the posterior distribution as $N_e^k \to \infty$. This follows from the usual importance sampling theory, see, e.g., [17], even if the proposal distribution depends on past samples (see also [15]). Moreover, at each iteration, a finite number of weighted samples are used to define a new proposal distribution, which we can chose such that its support includes the support of the posterior distribution. Thus, the next iteration is guaranteed to converge in the same sense. In summary, the iteration converges within one step, provided one generates sufficiently many samples,

i.e., if N_e^k is large enough. Since we want to use ISA on massively parallel computers, a large N_e at each iteration is feasible. Moreover, the iteration of ISA is serial, i.e., parallelism can only be leveraged by large N_e at each iteration. We thus wish to use ISA with a small number of iterations and large N_e at each iteration.

The importance of a large N_e^k at each step can be illustrated under simplifying assumptions. Suppose that the posterior distribution is Gaussian with variance Σ, and that the proposal distribution at iteration $k+1$ is also Gaussian. We neglect error in the mean, i.e., we assume that the mean of the proposal distribution is exactly equal to the mean of the posterior distribution, which can be justified because small errors in the mean are less severe than small errors in the covariances [60]. We further assume that the covariance matrix, computed from the weighted samples from the kth iteration, is $\hat{\Sigma} = (1+\varepsilon)\Sigma$, where ε is a small number. Under these assumptions, the weights at iteration $k+1$ are

$$w^{k+1} \propto \exp\left(-\varepsilon \, \frac{1}{2} x^t \Sigma^{-1} x\right), \quad x =\sim \mathcal{N}(0, \hat{\Sigma}),$$

which leads to

$$R^{k+1} = \left(\frac{1+\epsilon}{\sqrt{1+2\varepsilon}}\right)^{n_\theta} \approx 1 + \varepsilon^2 \cdot \left(\frac{n_\theta}{2}\right),$$

where n_θ is the number of parameters we wish to estimate, i.e., the dimension of θ. We can interpret ε as "sampling" error, and assume ε is inversely proportional to the square root of the effective sample size N_{eff} in equation (2). Thus,

$$R^{k+1} \approx 1 + R^k \, \frac{C \cdot n_\theta}{2N_e^k},$$

where C is a constant. The above equation indicates that the efficiency of the iteration is directly controlled by how many samples N_e^k one can produce at each step. However, the expansions above are only valid if ε is "small", which means that sampling error is small and the algorithm is already performing well with small R^k. However, studying the effects of *large* sampling error, and *small* ensemble size N_e^k on ISA is precisely what is needed in practice. Since theoretical considerations usually rely on very large ensemble sizes N_e^k, we argue that practical issues should be studied in the context of examples. We do not claim that ISA is efficient under broad conditions, but rather study applicability of ISA in two practical problems numerically, demonstrating that ISA can be efficient on several problems for which modern MCMC and importance sampling methods are slow to converge.

Finally, we wish to emphasize that we cannot guarantee that R decreases monotonically during the iteration when N_e^k is finite, in part because we must approximate R in equation (2) by averaging, and, thus, know R only up to sampling error, which can be substantial if R is large and N_e^k is not. Indeed, a large R is always uncertain, since, for example, an approximation of R by N_e samples is always less than N_e, however, the "true" R may be much larger. Since the iteration is stopped when the change in R from one iteration to the next is small, one must consider sampling error when choosing the tolerances.

3.2 Related literature on adaptive importance sampling and population Monte Carlo

Importance sampling algorithms that update the proposal distribution iteratively have been described in the literature before and are called "adaptive importance sampling", a related construction is called "population Monte Carlo". The earliest account we could find is the applications

paper by Buchner [12], where iterative sampling is used to compute failure probabilities in structural dynamics. The proposal distribution at each step is Gaussian and the mean and covariance of the Gaussian are updated at each step of the iteration. The paper also presents an initialization strategy for the iteration that works well in the application at hand. A more general version of ISA was described in [57, 58], where it is called adaptive importance sampling. Several stopping rules, similar to our convergence criterion based on R are presented, and consistency of the method is established, in the sense that expected values computed by ISA are asymptotically equal to expected values computed with respect to the posterior distribution. A method closely related to ISA is "population Monte Carlo" [15, 40], or "non-parametric importance sampling" [80]. In these methods, the ensemble at each step undergoes Kernel smoothing, i.e., the proposal distribution at iteration k is

$$q^k = \sum_{j=1}^{N_e} \alpha_j^k K_j(\theta, \theta_j^k),$$

where the sum of the α_j is one. Here, one can adapt the weights of the mixture α_j^k [27], or adapt the weights and the Kernels [14, 31] in each iteration. A recent application of population Monte Carlo in geosciences is presented in [74]. The scaling of the computational requirements of one step of population Monte Carlo with dimension of the parameter vector are discussed in a simplified setting in [42] (see also our conclusions in section 6).

We wish to use iterative algorithms for large scale applications using massively parallel computers. Parallelism allows us to draw a large number of samples at each step in the iteration, in which case Kernel smoothing becomes difficult. One can, of course, use a smaller number of mixtures at each iteration and adapt the weights of the mixtures and the Kernels, however here we use only *one* component, because this approach is sufficient to solve the parameter estimation problems we set out to solve.

3.3 Implementation and initialization of ISA

Implementation of ISA requires that one chooses the proposal distributions q^k. In principle, one could choose different classes of proposal distributions at each step in the iteration, e.g., a log-normal distribution during the first step, and a uniform distribution at the next step and so on. However, here we consider only the case where the class of proposal distributions is fixed to either Gaussians or multivariate t-distributions. For both distributions ISA only refines the mean and covariances during iteration (we keep the shape parameter ν of the multivariate t-distributions fixed). Thus, we first decide whether we wish to use Gaussians or multivariate t-distributions. Then, following the pseudo-code in algorithm 1, we compute mean and covariance from the initial ensemble to construct the initial proposal density. We then perform importance sampling, and, after weighting the samples, we re-compute mean and covariance to define the next proposal distribution. The utility of multivariate t-proposal distributions is discussed in a general setting in [60], where it is shown that multivariate t-proposal distributions can be particularly useful if one suspects that the posterior distribution has heavy tails.

One must also pick an ensemble size N_e^k at each step of the distribution. We advocate to chose this number as large as possible, since, under idealized conditions, one iteration is sufficient if N_e^1 is large enough, and since more samples imply that more of the parameter space is explored. Moreover, the accuracy of the estimation of R depends on how many samples we can generate. Indeed, a large N_e^k at each step is feasible because we can leverage parallelism in each iteration, which is the main advantage and motivation for developing ISA.

However, the assumption that a large number of samples can be generated does restrict the classes of problems we can attempt to solve with ISA, even when large parallel computers are used. For example, if a super-computer is needed to evaluate the model $\mathcal{M}$ once, then ISA is not feasible. Perhaps it is fair to state that ISA may enable parameter estimation using massively parallel computers, if the forward problem, i.e., evaluation of the model $\mathcal{M}$, can be done on a computer one scale below the computer we wish to use for ISA.

Initialization of ISA has can have a significant impact on the success of the algorithm. Indeed, the iteration cannot proceed if the initial proposal distribution q^0 yields only one effective sample, so that the variance that defines q^1 is zero, i.e., the algorithm has "collapsed" during the first step. We suggest two strategies for initializing ISA. The first is to run an MCMC code to obtain information about the mean and covariance of the posterior distribution. However, one has already solved the problem, and refinement of a proposal distribution by iteration is unnecessary, if the MCMC generates sufficiently many samples. We thus consider "short" MCMC runs, and in the examples and applications below this yielded good results. Moreover, the initialization strategy combines the robustness of MCMC with the computational advantages of importance sampling. In this context, one can think of ISA as a means to "fill out" an under-resolved MCMC computation, while making effective use of massively parallel machines. This can yield dramatic reductions in run times, from weeks or months of serial MCMC on a workstation, to minutes on an HPC system.

The second initialization strategy we test is to construct a Gaussian mixture model (GMM), for which each mixture component is centered at a mode of the posterior distribution. The various modes can be found in parallel, by initializing a suitable optimization code at various locations in parameter space obtained by, e.g., sampling the prior distribution. The covariances of the mixture are approximate Hessians of the negative log posterior distribution $F(\theta)$ at the modes. The weights of the mixture components are

$$\psi_j = \frac{\exp(-\phi_j)}{\sum_{i=1}^{n} \exp(-\phi_i)}, \tag{3}$$

where n is the number of minima we found and ϕ_i, $i = 1, \ldots, n$ are the local minima of F.

In this context one may find that the optimizer, initialized at two different points in parameter space finds the same minimum. We distinguish minima by considering the distance

$$d_{i,j} = (\mu_i - \mu_j)^T H_i^{-1} (\mu_i - \mu_j),$$

where μ_k are the various minima we find, and where H_k are the corresponding approximate Hessians with elements $(H_k)_{ij} = \partial^2 F / \partial\theta_i \partial\theta_j|_{\mu_k}$, $i, j = 1, \ldots, n$. We say that a minimum of F at μ_i is different from the one at μ_j if $d_{i,j}$ is more than a threshold. Other, more sophisticated assessments may be needed in other applications to account for asymmetries, since

$$(\mu_i - \mu_j)^t H_i^{-1} (\mu_i - \mu_j) \neq (\mu_i - \mu_j)^t H_j^{-1} (\mu_i - \mu_j),$$

However, in the applications we consider below, our simple strategy is sufficient.

3.4 Illustration by a 2D toy problem

We illustrate ISA by a simple "toy" problem, where we estimate a two-dimensional parameter vector θ. The prior is uniform on the cube $[0, 11] \times [0, 11]$, and the likelihood is such that the posterior is given by

$$p(\theta|z) \propto \exp(-F(\theta)), \quad \text{where } F(\theta) = 10^{-2} ||\theta - [5; 5]||^4 + 0.2 \sin(5||\theta||),$$

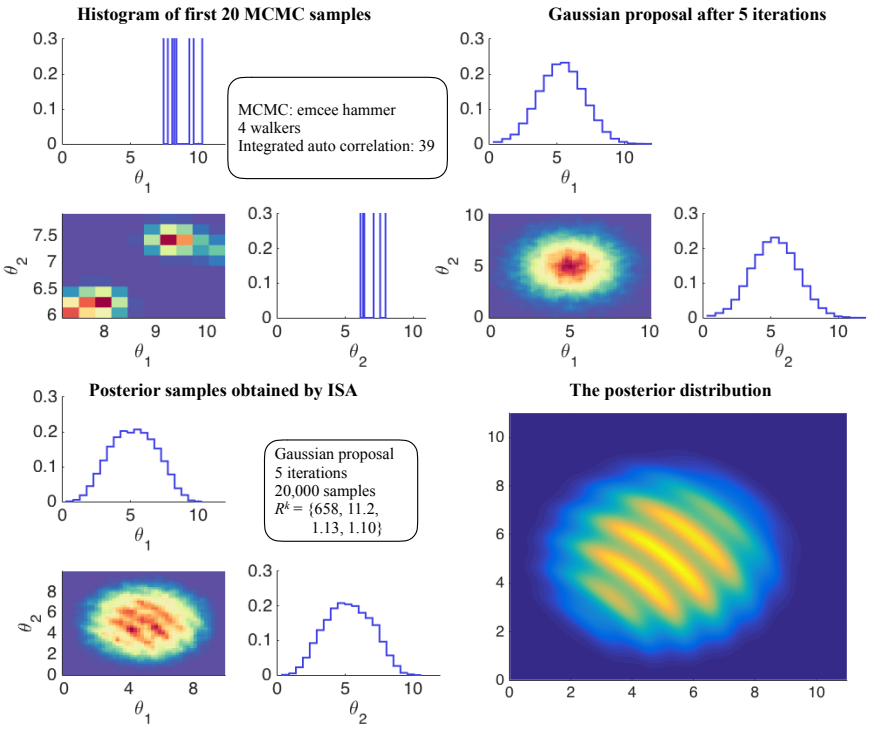

Figure 3: Illustration of ISA initialized by a short MCMC run. Top row, left: triangle plot of the first 20 MCMC samples. Top row, right: triangle plot of the Gaussian proposal distribution obtained after 5 iterations. Bottom row, left: triangle plot of the posterior distribution obtained by ISA. Bottom row, right: the posterior distribution, bird view.

and where $||x|| = \sqrt{x^T x}$ is the 2-norm. The posterior distribution is illustrated in the bottom right panel of figure 3. There are several modes that carry significant probability mass.

We apply ISA with Gaussian proposal distributions to obtain samples from this posterior distribution. We initialize the iteration by samples we obtain by the emcee implementation [30] of the affine invariant ensemble sampler [35]. We use four walkers, initialized by drawing samples from the uniform prior distribution. The first 20 samples are used to initialize ISA. A "triangle plot" of these 20 samples, which shows histograms of all one- and two-dimensional marginals of a given set of samples, is shown in the upper left panel of figure 3. We perform 5 iterations, each with $N_e = 20,000$ samples. The quality measure R converges quickly to about $R \approx 1.1$, by the sequence $\{R^k\} = \{658, 11.2, 1.13, 1.10\}$. A triangle plot from the Gaussian proposal we obtained after five iterations is shown in the upper right panel of figure 3. A triangle plot of 20,000 weighted samples of this Gaussian proposal are shown on the lower left panel of the figure, and one can clearly identify multiple modes in the posterior distribution. Moreover, the efficiency of importance sampling is almost perfect in this example: since $R \approx 1$, almost all samples are effective samples. Emcee hammer on the other hand is characterized by an integrated auto correlation time of about 39, which means that only one in 39 samples is effective. We estimated the integrated auto correlation time from an emcee run with four walkers, each taking 100,000 steps.

The convergence of ISA is further illustrated in figure 4, where we show histograms of one-dimensional marginals of the proposal distributions q^0, q^1, q^2, q^3, obtained during the iteration, as well as a histogram of the corresponding marginal of the posterior distribution. We observe that the high probability region of the initial proposal distribution is far from that of the posterior

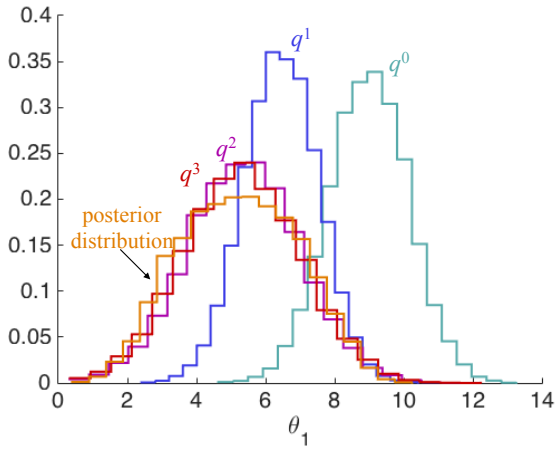

Figure 4: Convergence of the proposal distributions of ISA.

distribution. However, ISA quickly corrects this during iteration and moves the Gaussian proposal distributions closer to the posterior distribution.

We also tested ISA initialized by a GMM, which we obtain by performing 100 minimizations of F, each initialized by a sample from the uniform prior density. We then compare the minima by the distance $d_{i,j}$ as explained above, and find five distinct minima of F. At each distinct minimum, we approximate the Hessian by finite differences (we chose computationally more efficient methods to approximate Hessians in the applications below). The distinct minimizers and corresponding Hessians define the GMM with five mixture components, where the weights of each component is given by (3). We use this GMM as the proposal distribution and compute 20,000 samples to obtain $R \approx 2.59$. Using the GMM mixture distribution is thus less effective than using a few iterations of ISA as above. Nonetheless, we can use a few samples of the GMM to initialize ISA. Here we use the 50 samples from the GMM to initialize ISA and then iterate using Gaussian proposal distributions with 20,000 samples at each step. The Gaussian proposal distribution we obtain after three iterations is characterized by a quality measure $R = 1.10$ (the sequence is $R = \{3.27, 1.1, 1.1\}$). The experiments indicate the iterations of ISA yield similar proposal distributions after a few iterations for both initialization strategies.

4 Application 1: subsurface flow

4.1 Background and problem set-up

Subsurface flow models are extensively used in hydrological and geosciences applications, such as the prediction of the contaminant fate during a chemical spill, and the estimation of oil production from an oil recovery process. Robust predictions of quantities of interest require accurate characterization of the model parameters, such as the material properties and process-related parameters. Difficulties in direct measurements of these parameters necessitate the use of inverse modeling tools to infer them based on sparse measurements of observables that can be modeled.

In this section, we use ISA in a synthetic subsurface flow problem to estimate the permeability distribution, and the associated uncertainty, in a vadose zone from sparse measurements of saturations and change in the water volume in a domain that covers a two-dimensional 4m×3m (width×depth) spatial domain and is discretized by a mesh with uniform grid cells of size 0.05 × 0.05m. Water is released from a pond into a heterogeneous vadose zone that has a water table at

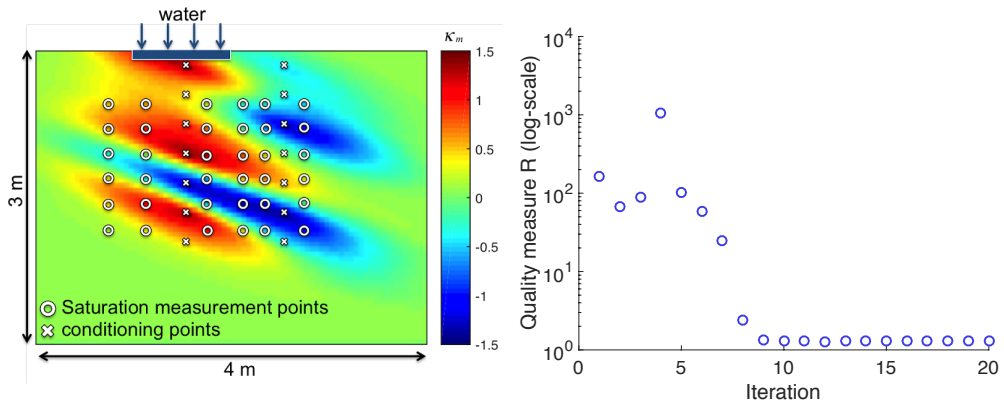

Figure 5: Left: Computational domain and locations of the conditioning points for Kriging; the color map illustrates the distribution of the permeability modifier when the values at the pilot points are given by θ_{ref}. Right: quality measure R during ISA iteration.

depth 3m. The saturation of the heterogeneous soil is initially in capillary-gravity equilibrium. A one-day ponded infiltration period is followed by a one-day water redistribution period. During infiltration, the water level in the pond is maintained at 2cm, and saturations at 36 locations, shown in the left panel of figure 5, are measured every two hours (25 snapshots) during the two days of infiltration and redistribution. The total volume of water flowing out of the pond is also measured at the same frequency, yielding a total of 925 observations that we combine into a vector z as in equation (1).

Our goal is to determine the permeability distribution of the vadose zone based on the above measurements. We describe the permeability by the pilot-point method [33, 64] with fourteen pilot points, located in two columns near the pond, as shown in the left panel of figure 5. The pilot points serve as adjustable conditioning points for determining the permeability modifier (κ_m) in each grid cell through a Kriging process that utilizes a spherical semivariogram model with fixed parameters [25]. The permeability in each grid cell is then given by $\kappa_{\text{ref}} \times 10^{\kappa_m}$, where $\kappa_{\text{ref}} = 0.316 \times 10^{-11}$ m^2 is a reference permeability value. The inverse problem we consider amounts to estimating the parameter vector $\theta = \{\kappa_{m,i}, 1 \leq i \leq 14\}$, where $\kappa_{m,i}$ are the κ_m values at the pilot point locations, based on two days worth of saturation and flow data.

We assume a broad Gaussian prior with zero mean and a diagonal covariance matrix C_{pp} where all 14 diagonal elements are equal to nine. The likelihood is defined by equation (1), where the model $\mathcal{M}$ is based on Richards' equation [67]. Numerically, we solve this equation using the integral finite-difference simulator TOUGH2 [63] with relative permeability and capillary pressure functions as in [77]. The errors v in equation (1) are Gaussian with mean zero and a diagonal covariance C_{zz}. The standard deviations that define the covariance are assumed to be 0.01 for the saturation measurements, and 0.005 m^3 for the water volume measurements. We perform experiments with synthetic data, i.e., we assume a reference permeability field defined by

$$\theta_{\text{ref}} = \{1.5, .5, 1., 1.5, -1.5, 1.5, 0.5, -0.5, -0.5, -1.5, -0.001, 0.5, -1.5, -0.5\},$$

and then perform a simulation of the above infiltration. The reference set of parameters defines the permeability modifier field shown in the left panel of figure 5. The synthetic observation vector z is then give by $\mathcal{M}(\theta_{\text{ref}}) + \epsilon$ where ϵ is an instance of v in equation (1).

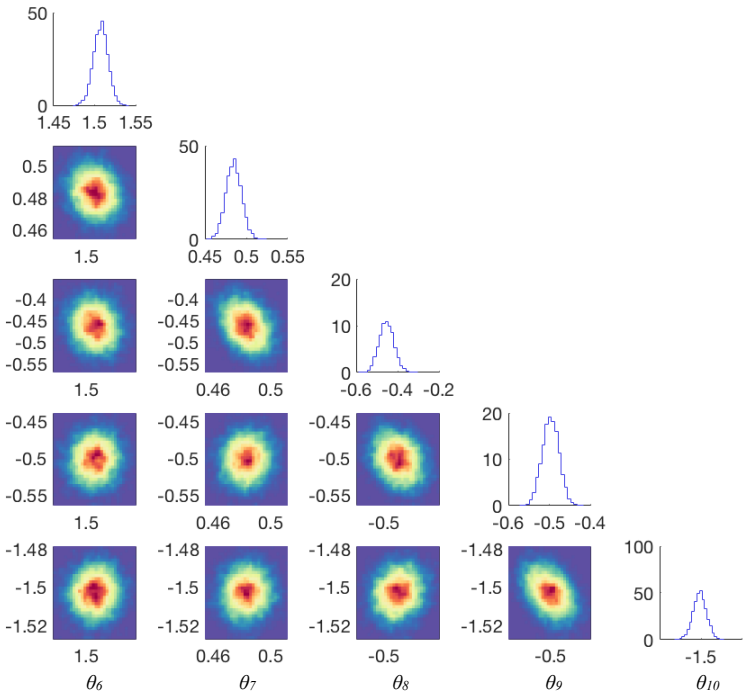

Figure 6: Triangle plot of posterior distribution obtained by ISA after nine iterations, each with 5,120 samples. The plot shows histograms of the parameters $\{\theta_6, \theta_7, \theta_8, \theta_9, \theta_{10}\}$ (diagonal), as well as two-dimensional histograms of all pairs of these five parameters (lower triangle).

4.2 Results

We initialize ISA by a single Gaussian. We define this Gaussian by solving the nonlinear least squares problem $\min_\theta F(\theta)$, where F is the negative logarithm of the posterior (up to an additive constant)

$$F(\theta) = \frac{1}{2}(z - \mathcal{M}(\theta))^T C_{zz}^{-1}(z - \mathcal{M}(\theta)) + \frac{1}{2}\theta^T C_{pp}^{-1}\theta. \tag{4}$$

Here $\mathcal{M}(\theta)$ is the output of the TOUGH2 simulator. Specifically, the minimizer of F defines the mean, and the approximate Hessian, evaluated at the minimum, defines the inverse of the covariance matrix. We approximat the Hessian by neglecting contributions from second derivatives, i.e., we set the covariance equal to $2\,J^T J$, where J is the Jacobian matrix of the nonlinear least squares problem. Numerically, we rely on Levenberg-Marquardt [43, 49] (lsqnonlin function in MATLAB) to solve the optimization problem and use finite differences for computing the first derivatives. For this particular problem, 20 iterations are needed to find a minimum. We note that the evaluations of the negative log-posterior (4) and of the Jacobian J in each iteration can be easily parallelized.

We start ISA using the Gaussian proposal distribution and generate 5,120 samples at each iteration. The likelihoods and weights of the samples are evaluated in parallel using 160 nodes, each with 32 cores, on Cori, a supercomputer at NERSC. With this set-up, the quality measure R converges from 163 to 1.3 in 9 iterations, as illustrated by the right panel of figure 5.

We compare ISA to solving the same problem by the affine invariant ensemble sampler [35], as implemented in the emcee code [30]. We use a parallel implementation of the algorithm with 126 walkers, which allows us to use 2 nodes on Cori efficiently. We discard the first 2,500 steps as "burn-in" and take another 2,500 steps per walker. The computational cost thus amounts to a

total of 630,000 likelihood evaluations, each requiring an evaluation of the TOUGH2 model that takes approximately 3 minutes. The integrated auto-correlation times for θ vary from 84 to 532, indicating that we have 3,750 to 592 "effective" samples. This should be compared to the quality measure of $R \approx 1.3$ after nine iterations of ISA. This means in particular that, after nine iterations, almost every Gaussian sample one generates contributes effectively to the numerical approximation of the posterior density. To get there, about 46,080 samples are used, and these can be evaluated in parallel in nine batches of 5,120 samples. Our results demonstrate that ISA can solve this parameter estimation problem as effectively as state-of-the-art MCMC, while achieving greater level of parallelism that reduces the wall-clock time. Including the cost of the initial minimization of (4), solving the problem by ISA takes about 90 minutes, while the emcee code runs for about 20 days.

Note that in the context of ISA, a quality measure R close to 1 can indicate that the posterior is nearly Gaussian. The reason is that we use a Gaussian proposal distribution, and $R \approx 1$ implies that the posterior distribution is also nearly Gaussian. This is illustrated by the triangle plot in figure 6. The plot shows histograms of the five parameters θ_6–θ_{10}, as well as two-dimensional histograms of pairwise combinations of these five parameters. Histograms of those parameters that are not shown in the figure are qualitatively similar. The figure suggests that the variables are indeed nearly Gaussian. While we can conclude, after the fact, that other methods specifically designed for nearly Gaussian problems (see literature review) may perform better on the present problem, the distributions of the parameters are not typically known *a priorily*. In the absence of prior knowledge on the parameters' probability distribution, our results show that ISA is generally effective, and in particular allows greater parallelism.

5 Application 2: combustion modeling

5.1 Background

Predictive combustion modeling entails solving a set of fluid conservation equations that are supplemented with models for thermodynamics, transport and chemical kinetics. For this study we assume an ideal gas mixture, gradient-driven transport, and a set of reversible Arrhenius rate expressions for chemical reactions. We focus on chemical combustion in syngas systems. Compared to fuels involving complex hydrocarbons or biofuels, the hydrogen-oxygen case is composed of a relatively small number of chemical species and reactions. For that reason, and perhaps since hydrogen-oxygen represents an important building block for more complex fuels, the basic form of the numerical model for hydrogen-oxygen systems, including the set of specific fundamental chemical reactions, is well-established. The parameters that define the model, however, remain open to considerable debate, as indicated by the number of models under development (e.g., [24, 32, 44, 51, 56, 69]). Note that the various parameter values reported over the years for each line of model development result from a long history of incremental optimization efforts that incorporate emerging experimental data sets. As a result, while different parameter sets may lead to similar predictive capabilities within a common range of their respective validated conditions, the models they define can generate dramatically different predictions outside that range (see the discussion, e.g., in [36], particularly related to Fig. 1, and in [26], related to Fig. 4).

The numerical model used here consists of conservation law PDEs, whose solution requires a database of parameters that characterize the relevant thermodynamic relationships, transport coefficients, and kinetic rate coefficients for the component gases. Our goal is to use ISA to incorporate measured experimental data into the characterization of uncertainty of a selected set of nine parameters that describe important behavior of this particular numerical model.

To a large extent, our study follows a recent effort to optimize 28 parameters of a hydrogen-oxygen combustion model [24]. In that work, the model was tuned specifically to predict measured values of 36 experiments, including constant-pressure flow reactors, constant-volume shock-tubes and steady, unstrained premixed flames. In the reference, a solution-mapping technique was employed to approximate the response of the model using a second-order polynomial. For each experiment, an a priori sensitivity analysis was used to reduce the set of active parameters in the response surface model. An objective function was constructed and minimized based on the square of model prediction errors, normalized by the uncertainty of the respective measurement. Prior knowledge of the parameters was used only to enforce bounds on the variation of each active parameter, and to inform the construction of the underlying polynomial response surface model. In addition, the work made a less direct use of prior knowledge in the form of the actual combustion model itself, including the list of all model parameters that are not tuned during optimization. In a Bayesian framework, the referenced optimization approach can be interpreted loosely as one based on a uniform prior over a cube whose sides are defined by arbitrary (prior-informed) parameter bounds. In the present work, discussed below, we estimate a subset of the parameters considered in [24] and quantify the respective uncertainties using ISA, based on the complete numerical combustion model, rather than solution-mapping. However, we make use of many of the same target experiments in order to quantify our likelihood estimate.

5.2 Prior and forward model, likelihood and data

We use the optimized mechanism of [24], shown in detail in Table 1, to define the mean of a truncated Gaussian prior distribution for the parameters shown in red. Since we wish to quantify the extent to which experimental data (see below) constrains the active parameters, we take the standard deviations of each parameter to be of the order of its mean value (as shown in the table). In this way, we generate a broad Gaussian distribution that extends well beyond the upper and lower bounds specified for each parameter (shown in the rightmost column of the table). We truncate the Gaussians at the bounds to obtain our truncated Gaussian prior distribution.

Our goal is to sharpen the prior knowledge by incorporating a limited set of experimental data. We consider the six flow reactors and eight ignition delay experiments labeled (*ign1a, ign1b, ign2a, ign2b, ign3a, ign4a, ign4b, ign5a, flw1a, flw2a, flw3a, flw4a, flw8a, flw8b*) in Table 1 of [24]. The data include Gaussian error estimates, taken as the standard deviations, σ, reported in Table 1 of [24]). We assume that the errors are uncorrelated, i.e., the covariance matrix that defines the errors v in equation (1) is diagonal. The data, assembled into a 14×1 vector z, and the Gaussian error model are used along with our numerical model to define a likelihood by equation (1).

In equation (1), each element of the vector z corresponds to one of the 14 experiments. We compute z using point-reactors subject to constant pressure (flow-reactor) or constant volume (shock-tube) constraints. VODE [11], is used to integrate either set of equations over time, and the solution is monitored until the desired diagnostic is obtained. Initial data for the simulations is provided in the description of the experiment in the reference, as is the event defining the measured quantity. For example, experiment *flw1a* reports the mean rate of change of H_2 mole fraction during the time interval that it is between 40 and 60 percent of its initial value. The computed and measured data, along with the sampled and prior mean values of the parameters, and the corresponding prior standard deviations and experimental measurement errors are combined for each sample vector to form the resulting likelihood.

Our combustion model is accessorized with several options for efficiency and robustness so that it can be driven by a number of sampling strategies. In particular, the model is allowed to fail gracefully if any component of a sample (a vector of nine parameters) is generated outside the

ID	Reaction	A*	β	E_a	σ	Bounds (L:U)
1	$H + O_2 = O + OH$	2.644(16)	-0.6707	17041	2(16)	1(16):5(16)
2	$O + H_2 = H + OH$	4.589(4)	2.7	6260	1(4)	1(4):1(5)
3	$OH + H_2 = H + H_2O$	1.734(8)	1.51	3430	2(8)	4(7):4(8)
4	$2\ OH = O + H_2O$	3.973(4)	2.4	-2110		
5	$2\ H + M = H_2 + M$	1.780(18)	-1	0		
6	$H + OH + M = H_2O + M$	4.400(22)	-2	0		
7	$O + H + M = OH + M$	9.428(18)	-1	0		
8	$2\ O + M = O_2 + M$	1.200(17)	-1	0		
9	$H + O_2\ (+M) = HO_2\ (+M)$					
	high pressure, $K_{f\infty}$	5.116(12)	0.44	0	5.116(12)	0:1.6(13)
	low pressure, K_{f0}	6.328(19)	-1.4	0	6.328(19)	0:2(20)
	TROE: $F_c = 0.5$					
	Third-body: O_2(0.85), H_2O(11.89), Ar(0.40)					
	He(0.46), H_2(0.75)					
10	$H_2 + O_2 = HO_2 + H$	5.916(5)	2.433	53502	5.916(5)	0:2(6)
11	$2\ OH\ (+M) = H_2O_2\ (+M)$					
	high pressure, $K_{f\infty}$	1.110(14)	-.37	0		
	low pressure, K_{f0}	2.010(17)	-.584	-2293		
	TROE: $F_c = 0.2654\exp(-T/94)$					
	$+0.7346\exp(-T/1756) + \exp(-5182/T)$					
	Third-body: H_2(6), H_2O(6), Ar(0.7), He(0.7)					
12	$HO_2 + H = O + H_2O$	3.970(12)	0	671		
13	$HO_2 + H = 2\ OH$	7.485(13)	0	295	7.485(13)	0:2(14)
14	$HO_2 + O = OH + O_2$	4(13)	0	0	4(13)	0:2(14)
15a	$HO_2 + OH = O_2 + H_2O$	2.375(13)	0	-500	2.375(13)	0:1(14)
15b	$HO_2 + OH = O_2 + H_2O$	1.000(16)	0	17330		
16a	$2\ HO_2 = O_2 + H_2O_2$	1.300(11)	0	-1630		
16b	$2\ HO_2 = O_2 + H_2O_2$	3.658(14)	0	12000		
17	$H_2O_2 + H = HO_2 + H_2$	6.050(6)	20	5200		
18	$H_2O_2 + H = OH + H_2O$	2.410(13)	0	3970		
19	$H_2O_2 + O = OH + HO_2$	9.630(6)	2	3970		
20a	$H_2O_2 + OH = HO_2 + H_2O$	2.000(12)	0	427		
20b	$H_2O_2 + OH = HO_2 + H_2O$	2.670(41)	-7	37600		

Table 1: Arrhenius rate parameters for H_2-O_2 combustion model (kinetics, and accompanying thermodynamics and transport parameter database taken from [24]). Parameters in red are active for the present study. The forward rate constant, $K_f = AT^\beta \exp(-E_a/RT)$. *The number in parentheses is the exponent of 10, i.e., $2.65(16) = 2.65 \times 10^{16}$. For the active parameters, standard deviation of the prior, σ and lower and upper bounds of the truncated Gaussian are provided in the last two columns.

bounds specified in Table 1. In this case, the likelihood is treated formally as zero. Similarly, some combinations of parameter values lead to predicted values that are very far from the corresponding measured data. Ignition delays, for example, may be prohibitively long to compute efficiently, or too short to accurately capture when scanning the solution at a "reasonable" sampling frequency. In either case, the diagnostic analysis routines can identify the failure mode and gracefully exit, signaling the driver code that the corresponding likelihood is to be taken as zero. Extra controls are provided/required in the specification of each experiment to help detect such scenarios. These auxiliary parameters are adjusted to have minimal effect on the presented data, but enable robust performance of the model evaluation over millions of sample sets. For each sample, the list of simulated experiments are evaluated in parallel over compute threads, using an OpenMP task manager. When a large set of independent sample vectors can be generated at once, the likelihood function itself can be evaluated in parallel over samples using MPI, where each MPI rank maintains its own version of the model parameter set and evaluates its likelihood using a local subset of the available cores.

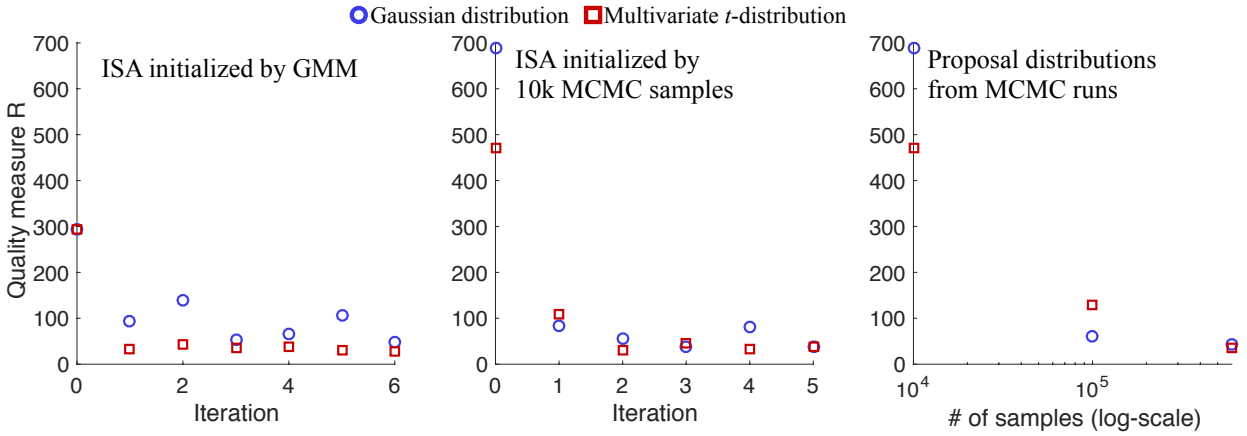

Figure 7: Iterations of ISA. Left: Quality measure R as a function of iteration number for ISA initialized by a GMM, using Gaussian or multivariate t-distributions and 10^5 samples at each iteration. Center: Quality measure R as a function of iteration number for ISA initialized by 10^4 samples of an MCMC run, using Gaussian or multivariate t-distributions and 10^5 samples at each iteration Right: Generating Gaussian/multivariate t-proposal distributions from 10^4, 10^5, and $6 \cdot 10^5$ samples.

5.3 Results

As explained in section 2.1, the prior distribution and a likelihood jointly define a posterior distribution, and we apply several variants of ISA algorithms to draw samples from this posterior distribution. We first test initialization by a GMM. To construct the GMM we run 10^4 minimizations of the negative logarithm of the posterior. The starting point for the optimizations are drawn from the broad prior. We use the nonlinear least squares code in the C implementation of "minpack" to carry out the required optimizations, and we use $2,400$ cores on NSERC's supercomputer "Edison". Of the 10^4 optimizations we tried, $2,482$ were "successful". The large number of failures can be explained by our broad prior, i.e., many of the starting points we tried were outside of the range the numerical model can simulate, which made the optimizer return a failure. We identified 12 minima to be distinct by using the criterion outlined in section 3.3. Here we use a threshold based on a χ^2 distribution so that the probability of the minima being within the threshold exceeds 95%. The resulting Gaussian mixture consists of 12 components, centered at the minima. The covariances are computed from approximate Hessians at these minima, using the same Gauss-Newton approximation as in section 4. We generate 10^5 samples from this GMM and evaluate the corresponding likelihoods using 3,072 cores on Edison. Evaluating and weighting 10^5 samples takes less than 15 minutes of wall-clock time, and results in the quality measure, $R = 293$. From the weighted samples we generate Gaussian or multivariate t-distributions with shape parameter $\nu = 3$ and iterate as described above. At each step of the iteration we "inflate" the sample covariance by a factor of two to mitigate sampling error and the underestimation of covariances, as is common in numerical weather prediction [1]. The inflation parameter could be tuned further, however we did not pursue this because we obtained good results. The left panel of figure 7 shows the quality measure R as a function of the iteration number. The figure suggests that ISA produces an effective proposal distribution after the first iteration with an associated R-value of less than 42. The precise sequence during iteration is $R = \{293, 32, 42, 35, 38, 29, 27\}$.

Also shown in the left panel of figure 7 are results we obtain by Gaussian proposal distributions

at each step of ISA. We note that the values of R fluctuate more than when using multivariate t-distributions, and the lowest value of R we obtain is also larger than the one we find by using multivariate t-proposal distributions. The precise sequence is $R = \{293, 94, 140, 54, 66, 105, 47\}$. Nonetheless, the values of R we obtain, in combination with the massive parallelism we can leverage at each step of the iteration, make ISA with Gaussian proposal distributions effective for this problem.

In the center panel of figure 7, we show results we obtain by ISA initialized by multivariate t or Gaussian distributions computed from 10^4 samples of an MCMC run. As before we use the emcee implementation of the affine invariant ensemble sampler [30, 35]. We describe the details of our MCMC run below, as it also serves as a reference solution. In the context of ISA, we note that the initial Gaussian/multivariate t-distribution yields a significantly larger R than the GMM we constructed above. The iteration however can quickly reduce R to values of around $R = 42$ (Gaussian proposal), and $R = 35$ (multivariate t-proposal). As above, we use 10^5 samples for each iteration. Initializing ISA with "few" MCMC samples requires at least two iterations to obtain R-values below 50, whereas we could reach low values of R after only one iteration when we initialize ISA by a GMM. Constructing the GMM is computationally more expensive due to the many optimizations we need to run, however constructing the GMM can leverage parallelism effectively. We conclude that, for this problem and computational set-up, the initialization by GMM is the more effective strategy.

We compare the ISA iterations to constructing Gaussian and multivariate t-proposal distributions from longer MCMC runs. These tests will illustrate (i) if the iterations can indeed speed up the sampling, or if we could simply produce slightly longer MCMC runs to obtain similar results; and (ii) to check the convergence of ISA on this problem. In the right panel of figure 7 we show the values of R we compute for proposal distributions as a function of the number of MCMC samples we use to compute the mean and covariances of these distributions. Specifically, we use 10^4, 10^5, and $6 \cdot 10^5$ to obtain $R = \{689, 60, 42\}$ for Gaussian proposals and $R = \{471, 130, 34\}$ for multivariate t-proposals. We thus obtain similar values of R as when iterating by ISA, provided that sufficient MCMC samples are used to construct the proposal distribution. This suggests that ISA can indeed find Gaussian or multivariate t approximations of the posterior distribution that can be used to generate posterior samples with efficiency of $R \approx 40$. More specifically, after one or two iterations, ISA produces one effective sample for every 40 samples we try. Each step of the iteration ISA can leverage parallelism, which results in significant savings in wall-clock time compared to using serial MCMC.

For our MCMC runs, we use a serial implementation of the emcee code, but parallelize the likelihood evaluations across the various experiments (flow reactors and ignition delays). In this way, we could effectively utilize the eight cores on a workstation. Generating $6 \cdot 10^5$ samples then requires about one month. For comparison, initializing ISA by 10^4 MCMC samples requires about 10 hours of MCMC for generating the initial ensemble, then 15 mins per iteration using 3,072 cores. Initializing ISA by GMM further reduces the wall-clock time required to generate the initial ensemble. On the other hand the emcee code also can be parallelized to a limited extent, which would reduce run-times and improve overall performance. Nonetheless, it is difficult to achieve the same level of parallelism as ISA by emcee or other MCMC codes. Moreover, the "shape" of the posterior distribution is difficult for the emcee code, which results in relatively long integrated auto-correlation times of a few hundred samples. Precise estimates of autocorrelation times are difficult to come by, however here we are satisfied with the estimates provided by the emcee code (which differ significantly across the nine variables). An autocorrelation of a few hundred means in particular that the "efficiency" of the emcee code is about one effective sample for every few hundred we try. This should be compared to the quality measure $R \approx 42$ of ISA (after a few

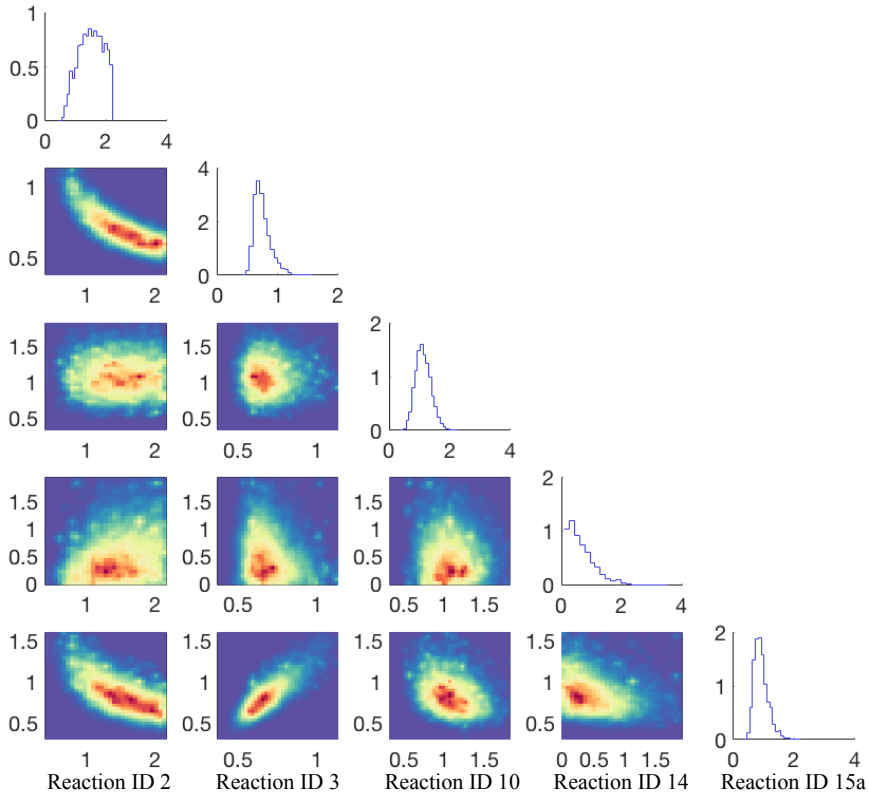

Figure 8: Triangle plot of posterior distribution obtained by ISA after six iterations, each with 10^5 samples. The plot shows histograms of reaction IDs $\{2, 3, 10, 14, 15a\}$ (diagonal), as well as two-dimensional histograms of all pairs of these five parameters (lower triangle).

iterations), which suggests that one in about 40 samples is effective.

Finally we show a triangle-plot of the posterior distribution we obtain by ISA after 6 iterations initialized by a GMM. The plot shows histograms of five parameters (Reactions IDs $\{2, 3, 10, 14, 15a\}$), as well as two-dimensional histograms of pairwise combinations of these five variables. In the plot, the samples are scaled by the prior mean as shown in table 1. We note that there is significant probability mass away from the prior mean, i.e., away from the value one. Moreover, the posterior mass is concentrated on narrow ridges in parameter space, which suggests that posterior variances are reduced significantly compared to the broad prior. We also note that there are rather strong correlations among some of the variables, e.g., reaction IDs two and three, and such correlations are absent in the prior. Thus, the experimental data we use here do indeed constrain the parameters we wish to estimate, and ISA can identify the parameter regimes that are consistent with the data we consider. We leave a more detailed discussion of implications of our study for combustion modeling for future studies, where we consider more complex models and more data, as our primary goal here is to demonstrate that ISA can be effective for sampling non-Gaussian posterior densities as they occur in many science problems.

6 Summary and conclusions

We presented iterative importance sampling schemes, which use samples obtained by importance sampling to construct a new proposal distribution that leads to more efficient importance sampling. These ideas have been discussed before (see our literature review), and the convergence of such iterations has also been studied. We revisited and revised these iterations in view of "realistic" test problems, and massively parallel computer architectures. In particular, we discussed and tested two strategies to initialize the iterations. One strategy uses a "short" MCMC run, while the other makes use of optimization and Gaussian mixture models.

It is important in practice that the iteration is robust when only few effective samples are produced at each step. We studied this behavior in two test problems, one in subsurface flow and one in combustion modeling. In our implementation of ISA we made use of massively parallel computers and found that the iterations can robustly leverage parallelism, which lead to significant speed ups when compared to MCMC.

A natural question is: how do the computational requirements of ISA scale with the dimension of the parameter vector? Our numerical results provide little information here because the dimension of the parameter vectors we considered are relatively small. On the other hand, we argued in section 3.1 that the quality measure R of Gaussian sampling problems depends exponentially on the dimension of the parameter vector (see also [16, 60]). This in turn implies that the number of samples required to achieve a given number of "effective samples" also increases exponentially with dimension. This may limit applicability of ISA, as the method would get stuck during the initialization phase if R were inevitably large in large dimensional problems. On the other hand, the Gaussian example that exhibits exponential dependence on the dimension does not account for, e.g., sparsity of covariance matrices of the target or proposal distributions. If such sparsity were present, and if one could exploit sparsity during sampling, then large dimensional problems may become within reach of ISA or other importance sampling methods. The situation is perhaps analogous to numerical linear algebra, where computations with large dense matrices are infeasible, while sparse matrices can be handled easily. In fact, sparsity of covariances is exploited during "covariance localization" of ensemble Kalman filters and is the key to being able to effectively solve estimation problems in numerical weather prediction (where the state dimension is several hundred million). If such strategies could be applied to parameter estimation problems, e.g., by enforcing that some data only inform a subset of the parameters one wants to estimate, then high-dimensional problems may become feasible. In the applications presented here, we enforce a "sparsity" by selecting, a priori, only a small set of parameters.

Acknowledgements

This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Applied Mathematics Program under contract DE-AC02-05CH11231, and by the Alfred P. Sloan Foundation through a Sloan Research Fellowship awarded to MM. This research used resources of the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

References

[1] J.L. Anderson. Spatially and temporally varying adaptive covariance inflation for ensemble filters. *Tellus*, 61(A):72–83, 2009.

[2] M.S. Arulampalam, S. Maskell, N. Gordon, and T. Clapp. A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking. *IEEE Transactions on Signal Processing*, 50(2):174 –188, 2002.

[3] E. Atkins, M. Morzfeld, and A.J. Chorin. Implicit particle methods and their connection with variational data assimilation. *Monthly Weather Review*, 141:1786–1803, 2013.

[4] J.M. Bardsley, A. Solonen, H. Haario, and M. Laine. Randomize-then-optimize: A method for sampling from posterior distributions in nonlinear inverse problems. *SIAM Journal on Scientific Computing*, 36(4):A1895–A1910, 2014.

[5] T. Bengtsson, P. Bickel, and B. Li. Curse of dimensionality revisited: the collapse of importance sampling in very large scale systems. *IMS Collections: Probability and Statistics: Essays in Honor of David A. Freedman*, 2:316–334, 2008.

[6] N. Bergman, editor. *Recursive Bayesian estimation: Navigation and tracking applications.* Ph.D Dissertation, Linkoping University, Linkoping, Sweden, 1999.

[7] A. Beskos, A. Jasra, E. Muzaffer, and A.M. Stuart. Sequential Monte Carlo methods for Bayesian elliptic inverse problems. *Stat. Comp.*, 25:727–737, 2015.

[8] P. Bickel, T. Bengtsson, and J. Anderson. Sharp failure rates for the bootstrap particle filter in high dimensions. *Pushing the Limits of Contemporary Statistics: Contributions in Honor of Jayanta K. Ghosh*, 3:318–329, 2008.

[9] C.J.F. Ter Braak. A Markov chain Monte Carlo version of the genetic algorithm Differential Evolution: easy Bayesian computing for real parameter spaces. *Statistics and Computing*, 16(3):239–249, 2006.

[10] C.J.F. Ter Braak and A.A. Vrugt. Differential evolution markov chain with snooker updater andfewer chains. *Statistics and Computing*, 18(4):435–446, 2008.

[11] P. Brown, G. Byrne, and A. Hindmarsh. Vode: A variable-coefficient ode solver. *SIAM Journal on Scientific and Statistical Computing*, 10(5):1038–1051, 1989.

[12] C.G. Bucher. Adaptive sampling – an iterative fast Monte Carlo procedure. *Structural Safety*, 5:119–126, 1988.

[13] T. Bui-Thanh, O. Ghattas, J. Martin, and G. Stadler. A computational framework for infinite-dimensional Bayesian inverse problems. Part i: The linearized case, with application to global seismic inversion. *SIAM Journal on Scientific Computing*, 36(4):A2494–A2523, 2013.

[14] O. Cappé, R. Douc, A. Guillin, J.-M. Marin, and C.P. Robert. Adaptive importance sampling in general mixture classes. *Statistics and Computing*, 18:447–459, 2008.

[15] O. Cappé, A. Guillin, J.-M. Marin, and C.P. Robert. Population Monte Carlo. *Journal of Computational and Graphical Statistics*, 13:907–929, 2004.

[16] A. J. Chorin, F. Lu, R.N. Miller, M. Morzfeld, and X. Tu. Sampling, feasibility, and priors in data assimilation. *Discrete and Continuous Dynamical Systems*, 36(8):4227–4246, 2016.

[17] A.J. Chorin and O.H. Hald. *Stochastic tools in mathematics and science*. Springer, third edition, 2013.

[18] A.J. Chorin and M. Morzfeld. Conditions for successful data assimilation. *Journal of Geophysical Research – Atmospheres*, 118:11,522–11,533, 2013.

[19] A.J. Chorin, M. Morzfeld, and X. Tu. Implicit particle filters for data assimilation. *Communications in Applied Mathematics and Computational Science*, 5(2):221–240, 2010.

[20] A.J. Chorin and X. Tu. Implicit sampling for particle filters. *Proceedings of the National Academy of Sciences*, 106(41):17249–17254, 2009.

[21] T. Cui, K. J. H. Law, and Y. M. Marzouk. Dimension-independent likelihood-informed MCMC. *Journal of Computational Physics*, 304:109–137, 2016.

[22] T. Cui, J. Martin, Y. M. Marzouk, A. Solonen, and A. Spantini. Likelihood-informed dimension reduction for nonlinear inverse problems. *Inverse Problems*, 29:114015, 2014.

[23] T. Cui, Y. M. Marzouk, and K. Willcox. Scalable posterior approximations for large-scale Bayesian inverse problems via likelihood-informed parameter and state reduction. *Journal of Computational Physics*, In press, 2016.

[24] S. G. Davis, A. V. Joshi, H. Wang, and F. N. Egolfopoulos. An optimized kinetic model for H_2/CO combustions. *Proc. Combust. Inst.*, 30:1283–1292, 2005.

[25] C. V. Deutsch and A. G. Journel. *GSLIB: Geostatistical software library and user's guide: Oxford Univ.* Oxford University Press, 1992.

[26] Y. Dong, A. T. Holley, M. G. Andac, F. N. Egolfopoulos, S. G. Davis, P. Middha, and H Wang. Extinction of premixed H2/air flames: Chemical kinetics and molecular diffusion effects. *Combus. Flame*, 142(4):374–387, 2005.

[27] R. Douc, A. Guillin, J.-M. Marin, and C.P. Robert. Convergence of adaptive mixtures of importance sampling schemes. *The Annals of Statistics*, 35(1):420–448, 2000.

[28] A. Doucet, N. de Freitas, and N. Gordon, editors. *Sequential Monte Carlo methods in practice*. Springer, 2001.

[29] H.P. Flath, L.C. Wilcox, V. Akçelik, J.Hill, B. van Bloemen Waander, and O. Ghattas. Fast algorithms for Bayesian uncertainty quantification in large-scale linear inverse problems based on low-rank partial Hessian approximations. *SIAM Journal on Scientific Computing*, 33(1):407–432, 2011.

[30] D. Foreman-Mackey, A. Conley, W. Meierjurgen Farr, D. W. Hogg, D. Long, P. Marshall, A. Price-Whelan, J. Sanders, and J. Zuntz. emcee: The MCMC Hammer. Astrophysics Source Code Library, March 2013.

[31] G.H. Givens and A.E. Raftery. Local adaptive importance sampling for multivariate densities with strong nonlinear relationships. *Journal of the American Statistical Association*, 91(433):132–141, 1996.

[32] P. Glarborg, M. U. Alzueta, K. Dam-Johansen, and J. A. Miller. Kinetic modelling of hydrocarbon/nitric oxide interactions in a flow reactor. *Combust. Flame*, 115:1–27, 1998.

[33] J. J. Gomez-Hernanez, A. Sahuquillo, and J.E. Capilla. Stochastic simulation of transmissivity fields conditional to both transmissivity and piezometric data–i. theory. *Journal of Hydrology (Amsterdam)*, 203(1):167–174, 1997.

[34] J. Goodman, K.K. Lin, and M. Morzfeld. Small-noise analysis and symmetrization of implicit Monte Carlo samplers. *Communications on Pure and Applied Mathematics*, page 10.1002/cpa.21592, 2015.

[35] J. Goodman and J. Weare. Ensemble samplers with affine invariance. *Communications in Applied Mathematics and Computational Science*, 5(1):65–80, 2010.

[36] J. F. Grcar, J. B. Bell, and M. S. Day. The soret effect in naturally propagating, premixed, lean, hydrogen-air flames. *Proc. Combust. Inst.*, 32:1173–1180, 2008.

[37] H. Haario, M. Laine, A. Mira, and E. Saksman. DRAM: Efficient adaptive MCMC. *Statistical Computing*, 16:339–354, 2006.

[38] H. Haario, E. Saksman, and J. Tamminen. An adaptive metropolis algorithm. *Bernoulli*, 7:223–242, 2001.

[39] K.M. Hanson. Markov Chain Monte Carlo posterior sampling with the Hamiltonian method. In *Proc. Sensitivity Analysis of Model Output*, 2001.

[40] Y. Iba. Population-based Monte Carlo algorithms. *Transactions of the Japanese Society of Artificial Intelligence*, 16(2):279286, 2000.

[41] M.H. Kalos and P.A. Whitlock. *Monte Carlo methods*, volume 1. John Wiley & Sons, 1 edition, 1986.

[42] J.E. Lee, R. McVinish, and K. Mengersen. Population Monte Carlo algorithm in high dimensions. *Methodology and Computing in Applied Probability*, 13:369–389, 2011.

[43] K. Levenberg. A method for the solution of certain problems in least squares. *Quarterly of Applied Mathematics*, 2:164–168, 1944.

[44] J. Li, Z. Zhao, A. Kazakov, and F.L. Dryer. An updated comprehensive kinetic model of hydrogen combustion. *Int. J. Chem. Kinet.*, 36:566–575, 2004.

[45] C. Lieberman, K. Willcox, and O. Ghattas. Parameter and state model reduction for large-scale statistical inverse problems. *SIAM Journal on Scientific Computing*, 32(5):2523–2542, 2010.

[46] J.S. Liu and R. Chen. Sequential Monte Carlo methods for dynamical systems. *Journal of the American Statistical Association*, 93:1032–1044, 1998.

[47] F. Lu, M. Morzfeld, X. Tu, and A.J. Chorin. Limitations of polynomial chaos expansions in the Bayesian solution of inverse problems. *Journal of Computational Physics*, 282:138–147, 2015.

[48] D.J.C. MacKay. Introduction to Monte Carlo methods. In M. I. Jordan, editor, *Learning in Graphical Models*, NATO Science Series, pages 175–204. Kluwer Academic Press, 1998.

[49] D.W. Marquardt. An algorithm for least-squares estimation of nonlinear parameters. *Journal of the Society for Industrial and Applied Mathematics*, 11(2):431–441, 1963.

[50] J. Martin, L. C. Wilcox, C. Burstedde, and O. Ghattas. A stochastic Newton MCMC method for large-scale statistical inverse problems with application to seismic inversion. *SIAM Journal on Scientific Computing*, 34:A1460–A1487, 2012.

[51] J. A. Miller. Chemical mechanism for methane combustion. privately distributed.

[52] M. Morzfeld, X. Tu, E. Atkins, and A.J. Chorin. A random map implementation of implicit filters. *Journal of Computational Physics*, 231:2049–2066, 2012.

[53] M. Morzfeld, X. Tu, J. Wilkening, and A.J. Chorin. Parameter estimation by implicit sampling. *Communications in Applied Mathematics and Computational Science*, 10:205–225, 2015.

[54] T. A. Moselhy and Y. M. Marzouk. Bayesian inference with optimal maps. *Journal of Computational Physics*, 231(23):7815–7850, 2012.

[55] G. Stadler N. Petra, J. Martin and O. Ghattas. A computational framework for infinite-dimensional Bayesian inverse problems: Part ii. stochastic Newton MCMC with application to ice sheet flow inverse problems. *SIAM Journal on Scientific Computing*, 36(4):A1525–1555, 2013.

[56] M. O'Connaire, H. J. Curran, J. M. Simmie, W. J. Pitz, and C. K. Westbrook. A comprehensive modeling study of hydrogen oxidation. *Int. J. Chem. Kinet.*, 36:603–622, 2004.

[57] M.-S. Oh and J.O. Berger. Adaptive importance sampling in Monte Carlo integration. Technical Report 89–19C, Department of Statistics, Purdue University, October 1989.

[58] M.-S. Oh and J.O. Berger. Adaptive importance sampling in Monte Carlo integration. *Journal of Statistical Computation and Simulation*, 41(3–4):143–168, 1992.

[59] D.S. Oliver. Minimization for conditional simulation: Relationship to optimal transport. *Journal of Computational Physics*, 265:1–15, 2012.

[60] A.B. Owen. *Monte Carlo Theory, Methods and Examples.* 2013.

[61] M. Parno and Y. M. Marzouk. Transport map accelerated Markov chain Monte Carlo. *Submitted*, 2016.

[62] M. Parno, T. Moselhy, and Y. M. Marzouk. A multiscale strategy for Bayesian inference using transport maps. *Submitted*, 2016.

[63] K. Pruess, G.J. Moridis, and C.M. Oldenburg. TOUGH2 User's Guide, Version 2.1, Report LBNL-43134. Report, Lawrence Berkeley National Laboratory, Berkeley, California, 1999.

[64] B.S. RamaRao, A.M. LaVenue, G. De Marsily, and M.G. Marietta. Pilot point methodology for automated calibration of an ensemble of conditionally simulated transmissivity fields: 1. theory and computational experiments. *Water Resources Research*, 31(3):475–493, 1995.

[65] P. Rebeschini and R. van Handel. Can local particle filters beat the curse of dimensionality? *The Annals of Applied Probability*, 25(5):2809–2866, 2015.

[66] N. Recca. A New Methodology for Importance Sampling. Master's thesis, Courant Institute of Mathematical Sciences, New York University, New York, USA, 2011.

[67] L.H. Richards. Capillary conduction of liquids through porous mediums. *Physics*, 1:318–333, 1931.

[68] S.L.Cotter, G.O.Roberts, A.M. Stuart, and D. White. MCMC methods for functions: modifying old algorithms to make them faster. *Statistical Science*, 28:424–446, 2013.

[69] G. P. Smith et al. GRI-Mech 3.0. Available at `http://www.me.berkeley.edu/gri_mech`.

[70] C. Snyder. Particle filters, the "optimal" proposal and high-dimensional systems. *Proceedings of the ECMWF Seminar on Data Assimilation for Atmosphere and Ocean.*, 2011.

[71] C. Snyder, T. Bengtsson, P. Bickel, and J. Anderson. Obstacles to high-dimensional particle filtering. *Monthly Weather Review*, 136(12):4629–4640, 2008.

[72] C. Snyder, T. Bengtsson, and M. Morzfeld. Performance bounds for particle filters using the optimal proposal. *Monthly Weather Review*, 143(11):4750–4761, 2015.

[73] A. Spantini, A. Solonen, T. Cui, J. Martin, L. Tenorio, and Y. M. Marzouk. Optimal low-rank approximations of Bayesian linear inverse problems. *SIAM Journal on Scientific Computing*, 37(6):A2451?A2487, 2015.

[74] A.S. Stordal. Iterative Bayesian inversion with gaussian mixtures: finite sample implementation and large sample asymptotics. *Computers in Geosciences*, 19:1–15, 2015.

[75] A. M. Stuart. Inverse problems: a Bayesian perspective. *Acta Numerica*, 19:451–559, 2010.

[76] A. Tarantola. *Inverse Problem Theory and Methods for Model Parameter Estimation*. Society for Industrial Mathematics, 2005.

[77] M.Th. van Genuchten. A closed-form equation for predicting the hydraulic conductivity of unsaturated soils. *Soil Science Society of America Journal*, 44(5):892–898, 1980.

[78] P.J. van Leeuwen, Y. Chen, and S. Reich. *Nonlinear Data Assimilation*. Springer, 2015.

[79] E. Vanden-Eijnden and J. Weare. Data assimilation in the low noise, accurate observation regime with application to the Kuroshio current. *Monthly Weather Review*, 2013.

[80] P. Zhang. Nonparametric importance sampling. *Journal of the American Statistical Association*, 91(435):1245–1253, 1996.